TECHNICAL WHITE PAPER

Rubrik and Microsoft Azure

Technology Overview and How It Works

TABLE OF CONTENTS

THE UNSTOPPABLE RISE OF CLOUD SERVICES

The digitization of business requires enterprises to move faster and be more agile to survive. Applying new technologies to existing business activities (e.g., how do we leverage AI to increase customer satisfaction?) fuels the cloud paradigm. According to IDC, more than 80% of IT organizations will be committed to hybrid cloud architectures by the end of 2017. For many enterprises, public cloud represents the ability to rapidly access resources for innovation while operating in a data-rich environment.

CLOUD PARADIGM INTRODUCES DIFFERENT PRINCIPLES

Enterprise IT looking to increase cloud usage will find that marrying non-cloud systems with cloud-native applications and infrastructure poses new principles.

- Shift from asset to service consumption. Traditional IT is largely based on providing finite assets that service relatively stable workloads and predictable business growth. In a cloud model, IT rapidly provisions services accordingly to business demand.
- Automate service delivery. With cloud, near-zero time to market can be delivered through automation frameworks. Infrastructure becomes programmable through code by being structured into templates that can be easily versioned and replicated for future deployments.
- Develop applications based on microservices. Rapid shifts in business demand require applications to deliver newer capabilities faster, to be resilient to failures, and to scale-out on-demand. Applications built in this new manner can be decomposed into independent components called “microservices”, each delivering a single function.

As public cloud plays a greater role in overall enterprise IT strategy, the need for a cloud-scale data management platform becomes paramount to protect and manage data born in the cloud and elsewhere.

WHAT IS RUBRIK CLOUD DATA MANAGEMENT?

As enterprises migrate applications to Azure, IT is depended on to deliver core data protection (backup, disaster recovery, archival) to maintain instant data accessibility through disasters, data loss, service outages, etc.

The Rubrik Cloud Data Management platform provides a cloud-native approach to managing the lifecycle of data, from creation to expiration, to drive better performance and operational continuity at lower costs. Rubrik bridges the gap between owned, on-premises infrastructure and the cloud by decoupling data from the data center through a software-defined fabric. Comprehensive data management is delivered through instant access, automated orchestration, and enterprise-class data protection and resiliency.

- Instant Access: Rubrik enables predictive global search and delivers instant application recovery by unifying data locked within disparate application silos into one globally indexed namespace while leveraging zero-byte cloning technology to enable on-demand copy data workflows.
- Automated Orchestration: Rubrik dramatically reduces daily operational management, providing a step-function change in simplicity by enabling a single policy engine to orchestrate service level agreements (SLAs) across the entire data lifecycle. The Rubrik programmatic interface automates how data services are created, consumed, and retired across data center and Azure.
- Security and Compliance: Rubrik secures data whether in-flight or at-rest throughout its lifecycle,

regardless of location. The Rubrik platform delivers granular role-based access control across all cloud data management workflows while providing automated compliance reporting to successfully complete various industry and internal audits.

DATA MANAGEMENT DESIGNED FOR AZURE

Delivering data protection and management for cloud requires a modern approach to accommodate the shift to service consumption, automation of service delivery, and development of modular, scale-out applications. Rubrik Cloud Data Management is designed with the following principles:

- Master-less, self-healing architecture: Rubrik distributes data, metadata, and task management across the cluster to deliver predictive scalability and eliminate performance bottlenecks. The system has its own distributed file system (Atlas) built from the ground up to store and manage versioned data at scale. Tasks are divvied up across cluster nodes based on data location and resource availability. Data is also stored efficiently while delivering resiliency (erasure coding).
- Distributed metadata and namespace: Rubrik's Distributed Metadata System operates alongside its cloud-scale file system (Atlas), providing a global index and catalogue that can be accessed at high speeds. It delivers continuous availability, linear scalability, and operational simplicity with no single point of failure in the cluster. The system is built to handle large amounts of data, distribute replicas of data across nodes (access to metadata is maintained even in the case of node failure), and provide low latency access to facilitate search.
- Policy-driven data management: Rubrik offers a global SLA policy engine in which users can automate protection of cloud applications, databases, and file sets to business requirements. Rubrik pioneers a declarative policy approach to eliminate the minutiae of scheduling data protection jobs–users simply select the desired snapshot frequency, retention duration, etc.
- Secure access in self-service environment: Granular control over user access is defined at a platform level, regardless of location. Rubrik allows self-service access (role-based access control) to empower users to perform their own backup, recovery, and archival services.
- Consumption and compliance transparency: Rubrik delivers actionable real-time operational insights for workload transparency, compliance, and capacity utilization across clouds (Rubrik Envision). Users can build their own dashboards, enable self-service access, and create custom reporting workflows with a responsive, HTML5-based interface and an API-first architecture.
- Easy integration with automation frameworks: Rubrik's API-first architecture enables automation of simple to complex data management workflows. Orchestrate custom workflows with a rich suite of RESTful APIs that can rapidly move local data to Azure, provision data management services for new cloud applications, and more.

HOW CLOUD DATA MANAGEMENT WORKS AND USE CASES

Rubrik Cloud Data Management is deployed on-premises through plug-and-play appliances and software appliances (Rubrik Edge). Alternatively, enterprises can deploy the Rubrik Cloud Data Management software on certified hardware platforms.

MANAGING CLOUD-NATIVE APPLICATIONS AND DATA

To protect cloud-native applications, Rubrik can be deployed as a software instance in Azure to orchestrate all critical data management functions – backup and recovery, replication and disaster recovery, archival, search, and more. Users spin up the recommended compute instance and scale easily by growing the Rubrik cloud cluster in lock-step with cloud data growth. All data is indexed and efficiently stored in a single, scale-out repository while providing data resiliency.

Rubrik provides the same consumer-grade HTML5 interface to manage the cloud as used to manage data on-premises and at the edge. Users can instantly locate (with real-time predictive search) and deliver application-consistent recoveries for data born in the cloud, including files, folders, file sets, VMs, and database instances (e.g., Windows, Linux, SQL databases). Users receive actionable insights with Rubrik Envision's rich visual reporting, which allows creation, customization, and sharing of platform analytics on consumption, compliance, and more, across a hybrid cloud environment.

By using Rubrik, enterprises free data from underlying infrastructure for ultimate workload portability. Keep in mind that workload portability across clouds can incur egress charges. Data transfer out from cloud to the internet will incur a charge. Data transfer within a public cloud service (data center to data center or region to region) will also often incur a charge.

Figure 1: Cloud Native Applications & Data

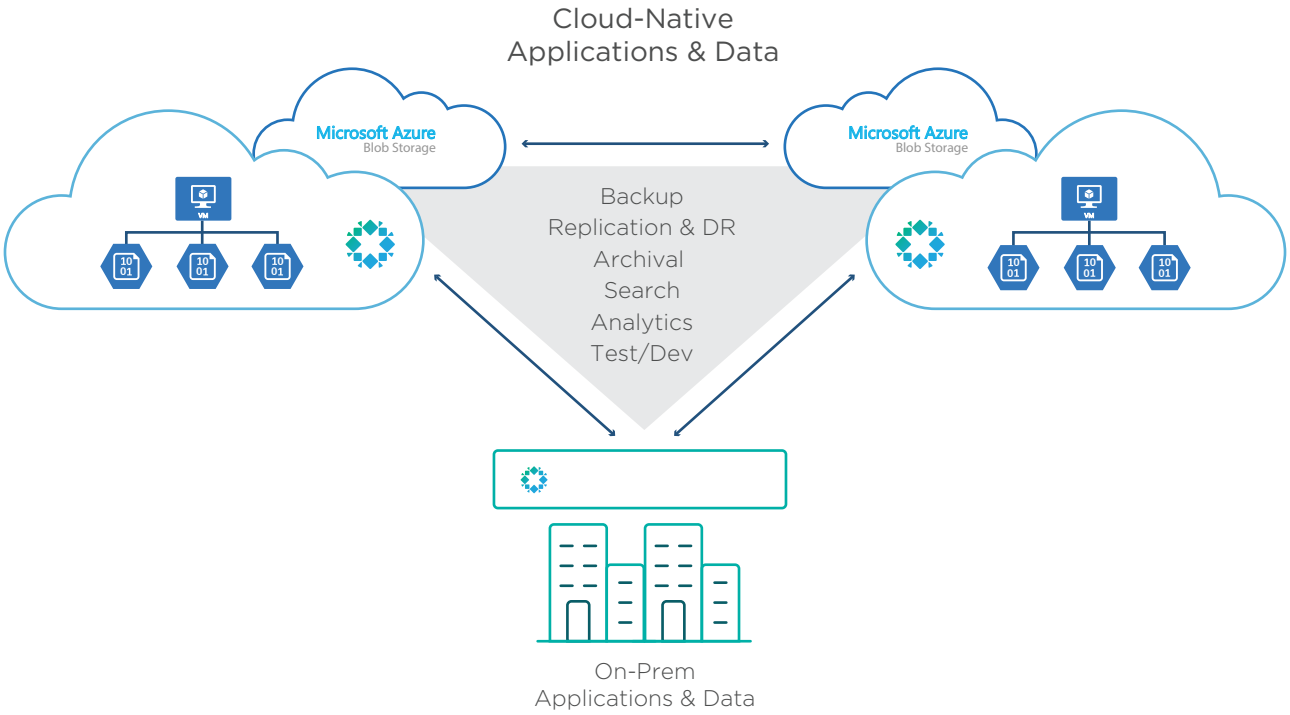

MANAGING HYBRID CLOUD APPLICATIONS AND DATA

Hybrid cloud enterprises can deploy Rubrik to manage applications on-premises while utilizing cloud services for archival, disaster recovery, and test/development.

Enterprises often ease into Azure by utilizing Blob Storage for long-term retention of application data while eliminating tape management complexity. With Rubrik, users can quickly set up archival policies to be applied across their hybrid cloud environment. Rubrik globally indexes all data, no matter where it resides, allowing users to retain quick access to archived data with predictive search. Users can instantly locate a file (rather than the entire data set) and download to restore in any location, saving both bandwidth and egress costs.

Enterprises can also leverage Rubrik to launch applications in Azure for application mobility, cloud disaster recovery, and test/development. To power on applications in Azure initially, users will configure the desired security group and Azure Virtual Network details. Rubrik scans the configuration file of a VM to understand its characteristics (compute, memory, storage, etc.) and recommends a compatible cloud instance type. At this point, Rubrik begins constructing a full snapshot in VHD from either data stored on-premises or in Azure Blob Storage. If the desired snapshot is present in Azure, a single ephemeral, lightweight Rubrik node is automatically created in the pre-configured Virtual Network to begin conversion of the VMware or Hyper-V VM into a VHD file to launch an Azure Virtual Machine in the cloud. This prevents the need for any data to exit the cloud region, saving both bandwidth and egress costs. Once the conversion is completed, the Rubrik node powers down and is purged until needed again.

Figure 2: App Instantiation in Microsoft Azure

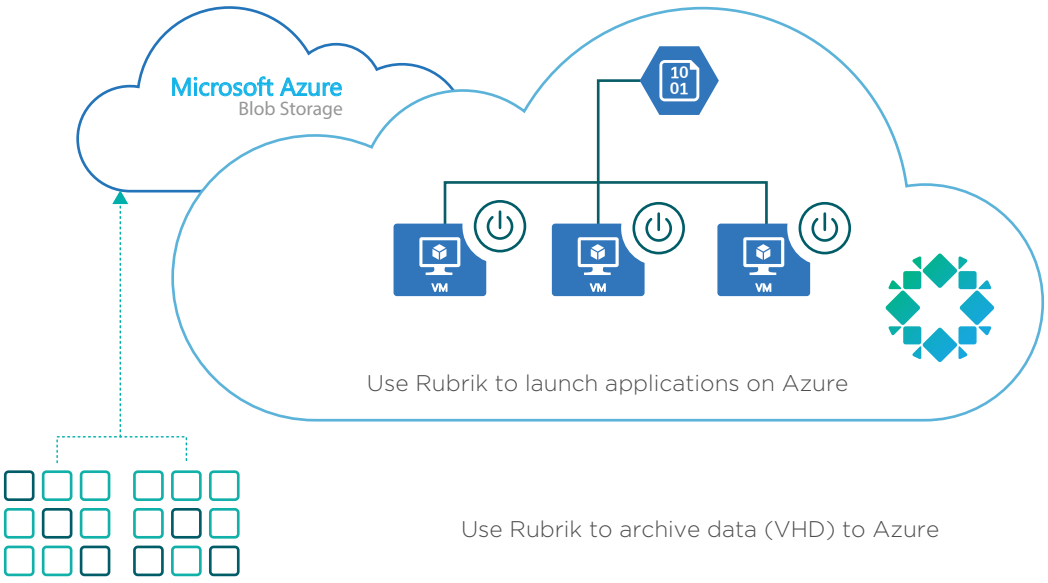

DATA MANAGEMENT FOR CLOUD-NATIVE APPLICATIONS	
Backup & Recovery of Cloud-Native Applications	Deploy Rubrik Cloud Data Management as a cloud instance and scale protection in-line with cloud service consumption. Protect cloud-native applications (such as Windows and Linux-based applications, SQL databases) by writing to Azure's storage services.
Cloud to Cloud Replication & Disaster Recovery	Deliver replication within a multi-cloud environment (different regions within Azure).
Cloud-Native Data Archival	Archive cloud-native data to Azure Blob Storage. Ensure instant accessibility of archived data with Rubrik's real-time predictive search.

DATA MANAGEMENT FOR HYBRID CLOUD APPLICATIONS	
Disaster Recovery to the Cloud	Instantiate your on-premises applications in Azure for disaster recovery. Rather than maintain a DR site complete with infrastructure that may largely remain idle, utilize cloud services when needed. Rubrik converts application data (VHD) into an Azure Virtual Machine. There is no need to run Rubrik in Azure to execute DR to the cloud, increasing overall cloud savings.
Migrating Test/Dev to the Cloud	Migrate existing on-premises applications to the cloud for test/development. Rubrik converts application data (VHD) into an Azure Virtual Machine. There is no need to run Rubrik in Azure to migrate workloads to the cloud for test/dev, increasing overall cloud savings.
On-Prem to Cloud Replication	Bi-directional replication is available from on-premises Rubrik cluster to Azure.
Data Archival	Send your application data to Azure Blob Storage for long-term retention while retaining immediate access with predictive search.

ENVIRONMENT SUPPORT CONSIDERATIONS	
Microsoft Azure	Standard D3 v2 instance Minimum of 4 instances (nodes), starting at 3 TB per node, for overall beginning usable capacity of 8 TB (erasure coding)

Rubrik converges backup software, replication, catalog databases, deduplicated storage, and more into a single software fabric. It does not retain the complex, bloated architecture reminiscent of legacy data protection solutions, whether it's running on-premises or in the cloud. The diagram below depicts the discrete software and hardware components eliminated by Rubrik through software convergence.

Whether Rubrik is deployed as a software instance in Azure or on-premises, users achieve multiple operational benefits, such as the ability to scale-out on a distributed architecture, to quickly automate data service workflows on an API-native framework, and to lessen infrastructure costs with a slimmer architecture.

Figure 3: Rubrik Eliminates Bloated, Legacy Architecture

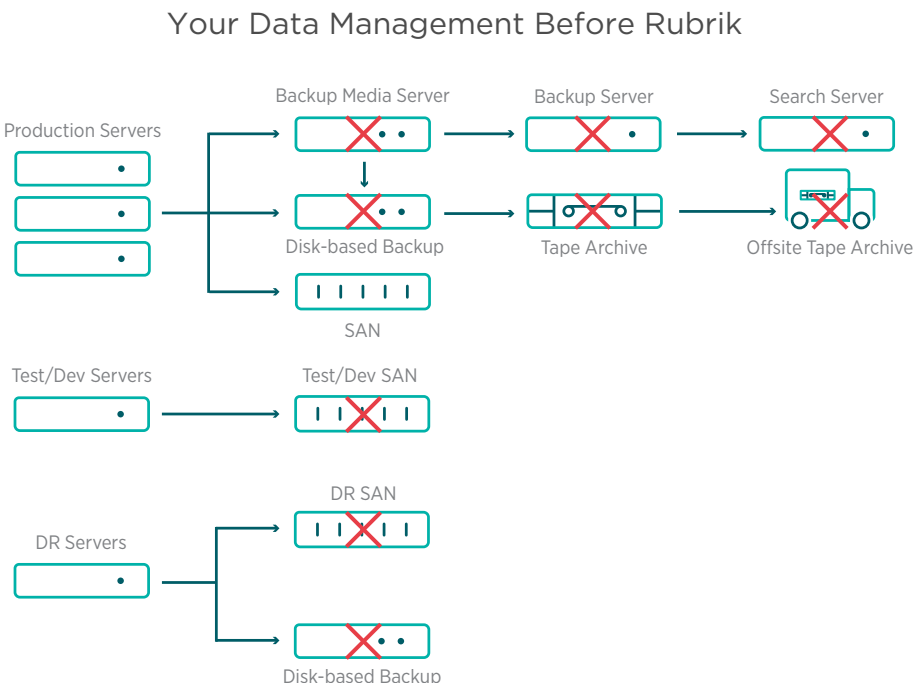

rubrik

Global HQ
1001 Page Mill Rd., Building 2
Palo Alto, CA 94304
United States

1-844-4RUBRIK
inquiries@rubrik.com
www.rubrik.com

Rubrik delivers instant application availability to hybrid cloud enterprises for recovery, search, cloud, and development. By using the market-leading Cloud Data Management platform to provide instant access with self-service, customers mobilize applications, automate protection policies, recover from Ransomware, search and analyze application data at scale on one platform. From days to seconds.

20170925_v1